Graph Anomaly Detection with Diffusion Model-Based Graph Enhancement (Student Abstract)

Shikang Pang[1], Chunjing Xiao[1], Wenxin Tai[2], Zhangtao Cheng[2], Fan Zhou[2]
[1]Henan University, Kaifen, Henan 475000, China
[2]University of Electronic Science and Technology of China, Chengdu, Sichuan 610054, China
pangsk0604@henu.edu.cn, chunjingxiao@gmail.com, wxtai@std.uestc.edu.cn, zhangtao.cheng@outlook.com, fan.zhou@uestc.edu.cn

Abstract

Graph anomaly detection has gained significant research interest across various domains. Due to the lack of labeled data, contrastive learning has been applied in detecting anomalies and various scales of contrastive strategies have been initiated. However, these methods might force two instances (e.g., node-level and subgraph-level representations) with different category labels to be consistent during model training, which can adversely impact the model robustness. To tackle this problem, we present a novel contrastive learning framework with the Diffusion model-based graph Enhancement module for Graph Anomaly Detection, DEGAD. In this framework, we design a diffusion model-based graph enhancement module to manipulate neighbors to generate enhanced graphs, which can efficiently alleviate the inconsistent problem. Further, based on the enhanced graphs, we present a multi-scale contrastive module to discriminate anomalies. Experimental results demonstrate the superiority of our model.

Introduction

As a few anomalies may cause tremendous loss, anomaly detection is a crucial task for a wide range of applications (Xiao et al. 2023b). Due to the difficulty of obtaining a large amount of annotated data, contrastive learning, which is in essence unsupervised learning, has been applied in graph anomaly detection (Liu et al. 2021). The core idea is to construct positive and negative pairs for contrast, following the principle of maximizing the agreement between positive pairs while minimizing that between negative pairs.

Since anomalies are frequently concealed across various scales (such as node and subgraph levels) due to the complexity of graph data (Jin et al. 2021), existing studies generally exploit node-level and subgraph-level representations to perform contrastiveness (Duan et al. 2023). They consider the node-level and subgraph-level representations derived from the same target node as a positive pair and the ones from the different target nodes as a negative pair. Despite the success of these ways, there still exists the additional issue: the two instances with different category labels might be considered as a positive pair, whose agreement is maximized during model training. This can adversely affect the model robustness and degrade the detection results.

Due to the severely skewed distribution of anomalies versus normal nodes in quantity, there exist a part of nodes who have heterophily-dominant neighbors, i.e., most neighbors have different class labels with the target node. Correspondingly, for a heterophilic node, its node-level representation derived from its attributes tends to have a different class label with its subgraph-level representation derived from its neighbors. However, general graph contrastive learning methods regard them as a positive pair and try to force them to be consistent (Jin et al. 2021; Liu et al. 2021). This can exert a negative impact on the model effectiveness.

To tackle these problems, we present a novel contrastive learning framework with Diffusion model-based graph Enhancement module for Graph Anomaly Detection, DEGAD. In this framework, we design a diffusion model-based graph enhancement method, which can manipulate neighbors to generate enhanced graphs to alleviate the problem of the inconsistent problem. Further, we adopt these enhanced graphs to construct a multi-scale contrastive module to detect anomalies.

Methodology

Problem Definition. Given an attributed network $G = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ with nodes $\mathcal{V}$, edges $\mathcal{E}$ and attributes $\mathbf{X}$, we aim to learn a function using a number of unlabeled samples to calculate the anomaly score for each node. The anomaly score represents the degree of abnormality.

Framework. There are two main modules in our proposed DEGAD. The graph enhancement module aims to manipulate the neighbors of the target node to produce enhanced graphs. The contrastive module leverages these enhanced graphs to train a GNN encoder with the designed multi-scale contrastive losses. As a result, the anomaly scores are computed based on the similarity degrees of the representations obtained from this model.

Graph Enhancement Module. Inspired by the superiority of diffusion models (Xiao et al. 2023a; Li et al. 2024), we present a diffusion model-based graph enhancement module, which can inject the characteristics of the target node into a part of its neighbors to generate enhanced graphs. The enhanced graphs can alleviate the inconsistent problem, because the subgraph-level representations can become more consistent with the node-level representation after injecting the characteristics of the target node into its neighbors.

This module first utilizes the forward process of the diffusion model to add noise into the source $\mathbf{z}_{\text{src}}$ to form a prior $\mathbf{z}^T$. Then, $\mathbf{z}^T$ is fed into the reverse diffusion process to generate a clean sample through gradual denoising, i.e., $\mathbf{z}^T \rightarrow \cdots \rightarrow \hat{\mathbf{z}}^t \rightarrow \mathbf{z}^t \rightarrow \cdots \rightarrow \mathbf{z}^0$. During this denoising process, the features of the reference nodes are extracted and iteratively injected into the latent variable $\hat{\mathbf{z}}^t$. In this way, the generated node $\mathbf{z}^0$ will possess the characteristics of both the source node and the reference nodes.

Concretely, to smoothly manipulate the neighbors, we select multiple nodes as the reference ones (e.g., $\mathbf{z}_{\text{ref}_1}...\mathbf{z}_{\text{ref}_n}$). Then, we utilize the forward process to generate a prior $\mathbf{z}^T$ by adding noise into the source $\mathbf{z}_{\text{src}}$:

$$\mathbf{z}^T = \sqrt{\bar{\alpha}_T}\mathbf{z}_{\text{src}} + \sqrt{1-\bar{\alpha}_T}\epsilon, \tag{1}$$

where $\bar{\alpha}_T = \prod_{t=1}^{T}(1-\beta_t)$ and $\beta_1, ..., \beta_T$ are fixed variance schedules. Further, we feed the prior $\mathbf{z}^T$ into the reverse diffusion process. In each denoising process, we exert the condition (the reference node) on this reverse process to generate a sample with the characteristics of both $\mathbf{z}_{\text{src}}$ and $\mathbf{z}_{\text{ref}}$:

$$\tilde{\mathbf{z}}^t = \gamma_1\hat{\mathbf{z}}^t + \gamma_2\big(\sigma\big(f_l(\mathbf{z}^t_{\text{ref}_1}), ..., f_l(\mathbf{z}^t_{\text{ref}_n})\big) - f_l(\hat{\mathbf{z}}^t)\big), \tag{2}$$

where $\mathbf{z}^{t-1}_{\text{ref}_1}$,..., $\mathbf{z}^{t-1}_{\text{ref}_n}$ are sampled using the forward process, $f_l(\cdot)$ is a low-pass filter, σ is a aggregation function like average, and γ_1 and γ_2 denote the weight hyper-parameters. This reverse process iteratively refines the distribution until reaching a clean manipulated sample.

Having the manipulated nodes, we replace a given ratio of the neighbors of the target node with these manipulated nodes to build an enhanced graph, called enhanced graph 1, and select another ratio for enhanced graph 2. These enhanced graphs are used for building the contrastive losses.

Contrastive Module. Based on the enhanced graphs, for a target node v_i, we first adopt a GCN to compute its node-level ($\mathbf{h}_i$) and subgraph-level ($\mathbf{e}_i$) representations. Regarding the node-level representation, its subgraph-level representation is considered as the positive pair, while the one corresponding to the other node is considered as the negative pair. In particular, for the enhanced graph 1, the anomaly level of the target node is correlated with the similarity:

$$S_i^1 = \varphi\left(\mathbf{h}_i^1\mathbf{W}_s\mathbf{e}_i^{1T}\right), \tag{3}$$

where $\mathbf{W}_s$ is a learnable matrix, and $\varphi(\cdot)$ denotes the Sigmoid activation. Generally, for similar degree S_i^1, the node-level and subgraph-level representations tend to be similar in positive pairs, i.e., $S_i^1 = 1$, and 0 otherwise. Hence, we employ the binary cross-entropy loss to train the contrast:

$$\mathcal{L}^1_{NS} = -\sum\big(y_i\log(S_i^1) + (1-y_i)\log(1-S_i^1)\big), \tag{4}$$

where y_i is equal to 1 in positive pairs and 0 in negative pairs. Similarly, for the enhanced graph 2, we can obtain the similarity degree S_i^2 and the loss $\mathcal{L}^2_{NS}$. Therefore, the final node-subgraph contrast loss is:

$$\mathcal{L}_{NS} = \lambda\mathcal{L}^1_{NS} + (1-\lambda)\mathcal{L}^2_{NS}, \tag{5}$$

where λ is a trade-off parameter for the two losses.

Method	Cora		CiteSeer		PubMed	
	F1	AUC	F1	AUC	F1	AUC
Autoencoder	0.6242	0.8703	0.5903	0.8045	0.6728	0.7880
DOMINANT	0.5062	0.8929	0.4627	0.8251	0.6917	0.8081
ALARM	0.6453	0.9124	0.6307	0.8431	0.7120	0.8257
CoLA	0.6683	0.8847	0.7102	0.8968	0.8412	0.9512
ANEMONE	0.6953	0.9122	0.7904	0.9189	0.8621	0.9548
GRADATE	0.7066	0.9237	0.8092	0.9409	0.9005	0.9820
DEGAD	**0.7101**	**0.9302**	**0.8241**	**0.9504**	**0.9153**	**0.9852**

Table 1: Performance comparison results.

Based on the similarity, we define the anomaly score: $AS_i = S_i^{\text{neg}} - S_i^{\text{pos}}$, where S_i^{neg} and S_i^{pos} represent the similarity degree of negative and positive pairs, respectively. The higher the scores, the more likely the nodes are anomalous.

Experiments

Datasets & Baselines. We conduct experiments on three datasets (Liu et al. 2021): Cora, Citeseer and Pubmed. We compare DEGAD with autoencoder-based models (e.g., Autoencoder (Zhou and Paffenroth 2017), DOMINANT (Ding et al. 2019) and ALARM (Peng et al. 2022)) and contrastive learning-based approaches (e.g., CoLA (Liu et al. 2021), ANEMONE (Jin et al. 2021) and GRADATE (Duan et al. 2023)).

Performance Comparison. We report the anomaly detection performance in Table 1 and have the following observations: (1) For overall detection results, DEGAD yields uniformly better performance than all the baselines across these datasets. The performance gain is primarily attributed to the enhanced graphs effectively advocating the quality of selected positive pairs. (2) Typical contrastive learning-based baselines outperform the non-contrastive learning models. This indicates that the contrastive learning-based pattern can effectively detect anomalies by mining the feature and structure information from graphs. While our module still surpasses these contrastive learning-based baselines.

Conclusion

In this paper, we presented a contrastive learning framework with the Diffusion model-based graph Enhancement module for Graph Anomaly Detection, DEGAD. In this framework, we designed a diffusion probabilistic model-based graph enhancement module to manipulate neighbors to generate enhanced graphs, which can relieve the problem of inconsistent positive pairs. Based on the enhanced graphs, we presented a multi-scale contrastive module to discriminate anomalies. The experiments demonstrate that our proposed approach achieves state-of-the-art performance.

Acknowledgments

This work was supported in part by National Natural Science Foundation of China (Grant No.62176043 and No.62072077).

References

Ding, K.; Li, J.; Bhanushali, R.; and Liu, H. 2019. Deep anomaly detection on attributed networks. In *SDM*, 594–602.

Duan, J.; Wang, S.; Zhang, P.; Zhu, E.; Hu, J.; Jin, H.; Liu, Y.; and Dong, Z. 2023. Graph Anomaly Detection via Multi-Scale Contrastive Learning Networks with Augmented View. In *AAAI*, 7459–7467.

Jin, M.; Liu, Y.; Zheng, Y.; Chi, L.; Li, Y.-F.; and Pan, S. 2021. Anemone: Graph anomaly detection with multi-scale contrastive learning. In *CIKM*, 3122–3126.

Li, X.; Xiao, C.; Feng, Z.; Pang, S.; Tai, W.; and Zhou, F. 2024. Controlled graph neural networks with denoising diffusion for anomaly detection. *Expert Systems with Applications*, 237: 121533.

Liu, Y.; Li, Z.; Pan, S.; Gong, C.; Zhou, C.; and Karypis, G. 2021. Anomaly detection on attributed networks via contrastive self-supervised learning. *TNNLS*, 33(6): 2378–2392.

Peng, Z.; Luo, M.; Li, J.; Xue, L.; and Zheng, Q. 2022. A deep multi-view framework for anomaly detection on attributed networks. *TDKE*, 34(6): 2539–2552.

Xiao, C.; Gou, Z.; Tai, W.; Zhang, K.; and Zhou, F. 2023a. Imputation-based Time-Series Anomaly Detection with Conditional Weight-Incremental Diffusion Models. In *KDD*, 2742–2751.

Xiao, C.; Xu, X.; Lei, Y.; Zhang, K.; Liu, S.; and Zhou, F. 2023b. Counterfactual Graph Learning for Anomaly Detection on Attributed Networks. *TKDE*, 35(10): 10540–10553.

Zhou, C.; and Paffenroth, R. C. 2017. Anomaly detection with robust deep autoencoders. In *KDD*, 665–674.